

BBN Report No. 2459


31 July 1972

AD 749378

Resource Sharing Computer Communication Networks

by

Robert E. Kahn



The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Advanced Research Projects Agency or the U.S. Government.

This research was supported by the Advanced Research Projects Agency under ARPA Order No. 1967; Contract No. DAHC-71-C-0083.

Distribution of this document is unlimited. It may be released to the Clearinghouse, Department of Commerce for sale to the general public.

NATIONAL TECHNICAL
INFORMATION SERVICE

DOCUMENT CONTROL DATA - R & D

Security classification of title, body of abstract and indexing annotation must be entered when the overall report is classified

1. ORIGINATING ACTIVITY (Corporate author) Bolt Beranek and Newman Inc. 50 Moulton Street Cambridge, Mass. 02138		2a. REPORT SECURITY CLASSIFICATION Unclassified	
		2b. GROUP	
3. REPORT TITLE RESOURCE SHARING COMPUTER COMMUNICATION NETWORKS			
4. DESCRIPTIVE NOTES (Type of report and, inclusive dates) Scientific			
5. AUTHOR(S) (First name, middle initial, last name) Robert E. Kahn			
6. REPORT DATE 31 July 1972		7a. TOTAL NO. OF PAGES 52 600	7b. NO. OF PAGES 34
8a. CONTRACT OR GRANT NO. DAHC15 71 C 0088		9a. ORIGINATOR'S REPORT NUMBER(S) BBN Report No. 2459	
b. PROJECT NO. ARPA ON 1967		9b. OTHER REPORT NO(S) (Any other numbers that may be assigned this report)	
c.			
d.			
10. DISTRIBUTION STATEMENT Distribution of this document is unlimited. It may be released to the Clearinghouse, Department of Commerce for sale to the general public.			
11. SUPPLEMENTARY NOTES This research was sponsored by the Advanced Research Projects Agency under ARPA Order No. 1967.		12. SPONSORING MILITARY ACTIVITY	
13. ABSTRACT <p>The development of resource sharing networks can facilitate the provision of a wide range of economic and reliable computer services. Computer Communication Networks allow the sharing of specialized computer resources such as data bases, programs and hardware. Such a network consists of both the computer resources and a communications system interconnecting them and allowing their full utilization to be achieved. In addition, a resource sharing network provides the means whereby increased cooperation and interaction can be achieved between individuals. This paper provides an introduction to computer networks and resource sharing and discusses some aspects of distributed computation.</p>			

1A

KEY WORDS

LINK A

LINK B

LINK C

ROLE

WT

ROLE

WT

ROLE

WT

resource sharing
computer communications
networks
time-sharing
Distributed Systems
data communications
message switching
circuit switching
IMP
Host Protocol

Resource Sharing Computer Communication Networks

by

Robert E. Kahn

Bolt Beranek and Newman Inc.

Cambridge, Mass. 02138

31 July 1972

This work was supported by the Advanced Research Projects Agency
of the Department of Defense under contract No. DARC-71-C-0088.

TABLE OF CONTENTS

I.	INTRODUCTION	1
II.	DISTRIBUTED OR CONCENTRATED RESOURCES.....	4
III.	COMPUTER TO COMPUTER COMMUNICATIONS	7
IV.	MESSAGE SWITCHED COMMUNICATIONS	11
V.	NETWORK USE OF INDIVIDUAL COMPUTER SYSTEMS ..	24
	a. Network Access	25
	b. User Requirements	27
	c. Message Processing	30
VI.	APPLICATIONS OF MULTI-COMPUTER INTERACTION ..	34
	i. File Transfer	35
	ii. Remote Job Services	36
	iii. Multi-Process Operation in Many Machines	38
VII.	A DISTRIBUTED OPERATING SYSTEM	41
VIII.	CONCLUSIONS	46
	REFERENCES	49

I. INTRODUCTION

The interaction between computers and communications has steadily developed over the last two decades. While many universities, government agencies and business firms prefer to make use of their own computers, an increasing number of people are using communication facilities to access commercial computer services [1]. Time-sharing and batch processing services are offered in most major U.S. cities or are accessible via telephone circuits, and communication charges for local telephone access to these services are, in general, substantially lower than the computer charges. As the use of computer services increases, the demand for reliable and low cost means of communicating over wide geographic areas also increases.

For many years, networks of interconnected computers have been planned or under study, and more recently several have been under development [16,18,24]. A common objective underlying the interest in these networks has been to demonstrate that economic savings or increased capabilities are possible by sharing computer or communication resources. Program access to specialized data bases is an important example of resource sharing in a computer to computer network.

The growing usage of these data processing services and the objective of sharing resources raises communication issues far more extensive than those of achieving increased capability and lower costs in the telephone network, or developing improved

communication services [2]. They involve a number of complex regulatory issues, the need for common methods of access to and interchange between data processing systems, the pooling of computer resources for increased utilization and reliability, the provision of specialized services, data conferencing and so forth. A set of associated regulatory issues involving telecommunications policy have been raised and are under intensive study. Are separate common carrier data networks desirable or not? What is the most effective way to plan for interconnection of data networks and how should their usage be tariffed?

It is too early to accurately predict in what way this interaction of computers and communications is likely to evolve. The technology is changing rapidly, and regulatory policies are in flux. If communication costs are not to dominate the overall cost of using remote data processing services, technological advances must allow communications at substantially lower per bit costs than are possible with the current switched telephone plant.

In this paper, we present one view of computer communications network development and explore a number of the important issues in distributed computation which have arisen. This paper is neither a completely general treatment of computer networks, nor a full case study, but rather it contains selected aspects of the two. The reader will, no doubt be able to

identify where general considerations give way to specific ones derived primarily from the author's experience with the development of the ARPANET* [7,18]. It is impossible for this to be an exhaustive treatise, or even a comprehensive one, and no such attempt is made.

*The Advanced Research Project Agency Computer Network.

II. Distributed or Concentrated Resources

Many economic factors support the conclusion that geographic "clustering" of computers is a desirable strategy for computer service organizations [9]. One possible advantage is better equipment utilization due to the pooling of resources. Clustering implies that a single maintenance staff (which is often underutilized) and scarce system personnel can support more equipment, more reliably, and that space, auxiliary equipment and overhead can be consolidated. In fact, several commercial time-sharing firms have already chosen to concentrate their computer resources in a small number of geographic areas. In contrast, however, many individual research or development machines under private or government ownership are distributed throughout the U.S. The valuable resources on many of these machines provide a strong incentive for them to be made available to users and computers at many other locations [10,11,33].

The location of computers at a few geographic locations requires that both local and remote users be provided with an economic and reliable way to access the service. The switched telephone network currently appears to be a poor candidate to provide the long distance communications service. In addition to being considerably more costly than local service, the error performance on long distance circuits is degraded from shorter circuits and is insufficient for many computer applications. In addition, frequent disconnections, busy signals, etc. during

peak traffic hours often make its usage inconvenient. These factors, coupled with user desires for increased bandwidth, lower set-up times, and more suitable tariffs have encouraged several vendors to competitively enter the common carrier market [28].

The tariffs and the technical characteristics of the circuit switched telephone network reflect the nature of voice communication requirements that are quite different from those of computer communications. Due to the bursty nature of computer traffic and the extremely low utilization of a typical voice grade circuit by a terminal, a substantial portion of the data communications capacity in a circuit switched system is simply not used. This results in inefficient utilization of telephone company resources from the users point of view. Frequency or time division multiplexing techniques have been usefully applied for deriving individual channels, but the statistical nature of computer traffic makes fixed allocation strategies such as these inefficient or unacceptable.

On the other hand, statistical multiplexing techniques allow these circuit resources to be more widely shared, at the possible expense of occasional delays in transmission. Message switching employs a generalized form of multiplexing for a network environment that allows all circuits to be shared among all users in a statistical fashion without being allocated in advance.

This has been the motivation for the development of new communication systems as well as combined computer communication networks. The construction of both common carrier data communication systems and "private" networks (using leased common carrier facilities) is a natural outcome of the need for economic and reliable communication between users and geographically distributed computers. In addition to potential cost savings, many of these networks provide error control, as well as asynchronous operation, local echoing, speed and code conversion, which are better suited to data communication with computers than use of the telephone network alone. A reevaluation of the tariff structures for data communication has recently been undertaken by the F.C.C. and efforts are being made to provide the public with data communications service having lower error rates, smaller service charging intervals, and faster set-up times than the switched voice network currently provides.

III. Computer to Computer Communications

A computer network is a complex collection of many types of resources, including data bases, programs, operating systems, and special purpose hardware, all of which are capable of being accessed from any other resource in the net. Computer to computer communication is necessary to achieve effective resource sharing, but the ability to transfer information between machines does not automatically result in useful machine to machine interactions. Aroused by the exciting possibilities in using multiple machines, system designers have recently begun to provide the major technical effort required to achieve effective computer to computer communication. The existence of the ARPANET is having precisely this effect, and as a result the extent of computer to computer interactions is certain to grow substantially in the next few years [6].

The ARPANET is one of the most advanced examples of a computer communication network [8,16,18]. It consists of a geographically distributed set of different computers, interconnected by a communication system based upon very fast response (interactive) message switching. This network was developed to ultimately allow economic and reliable sharing of specialized computer resources. The ARPANET has demonstrated the feasibility of message switching technology, illustrated its advantages, and fostered the development of techniques for computer to computer communication. It is interesting to note

that the ARPANET was originally designed with the notion of computer to computer communication in mind. It has subsequently been extended in capability to allow users with terminal equipment but no computer to connect to the net and communicate with computers and other users. In this sense, the ARPANET has taken the opposite approach from every other network designed with user access originally in mind.

For many years, the National Physical Laboratory (NPL) in England has experimented with the use of "single packet" messages for switching in the "local area of a data communication network" [31]. A number of terminal devices were successfully interconnected into a local network at NPL and recently they have been concerned with extending the local network into a distributed network [30]. A computer to computer network is also under development in France to allow data sharing without costly duplication of files, and its attendant problems of control, updating, security etc. Central files, each accessible via a local computer, will be made accessible to other computers and hence to an extended user community. This network is expected to use a message switching technique similar to that used in by the ARPANET in the U.S. In addition, networks are under design or development in other countries (e.g. Canada, Japan).

In general, the properties and structure of a computer communications network must reflect the overall requirements for which it was designed. This may consist of high speed

(megabit/second) circuits for rapid computer to computer communication, or low speed (voice and telegraph grade) circuits for terminal access or slow speed communications; it may be circuit switched or message switched, etc. Whatever its detailed structure the network contains a communication system (private or common carrier) and a set of computer system resources and users that interact via the communications system. This system is also called a communication subnet or simply a subnet for short. This organization not only characterizes the organization of geographically distributed networks, but can also serve as a model for the local structure of a single computer complex [3]. Its structure is therefore quite fundamental.

In operating a computer communications network as a "marketplace" for computer related services, a number of important issues arise [34]. We allude to a few of them here. What criteria are appropriate to determine whether a service may be removed from the system? When and where should additional services be incorporated and what procedures are needed to maintain effective competition? What subnet changes are appropriate for changes in the distribution of resources? The total operational procedure should also include a strategy for utilization of the resources consistent with its intended functions (e.g. load sharing, data sharing etc.)

An overriding concern of the network design is the overall reliability of the communications and computer resources. For a

user to entrust his computing to a network, he must develop confidence in its availability when he needs it. It must be convenient to use and it must provide a believable guarantee to maintain standard and expected grades of service. An investment in time and energy to use a network resource can be negated by the failure to maintain a consistent service offering. Insuring that proper concern exists for the remote user of a computer resource is an important administrative problem that affects almost every phase of computer network development.

IV. MESSAGE SWITCHED COMMUNICATIONS

Since the message switching technology is not as well established as the circuit switched technology, the fundamentals of its operation are reviewed in this section. Considerable discussion on the nature of these two switching doctrines is taking place. Are they merely different ends of a common spectrum (with a key variable such as packet size), or are they fundamentally different communication techniques? An argument in favor of their similarity is that both types rely on store and forwarding of data, whether a single bit is transiently stored, a byte sized envelope or a larger sized packet. The most significant external characteristics that "appear" to distinguish the two systems are that 1) circuit switching systems are better equipped to maintain a time frame for users that require continuity in transmission, as in speech, while 2) message switching systems allow speed and code conversion, thus permitting direct connection of and communication between devices of widely varying type. But it is possible to mask even these "seemingly" essential differences by the provision of a small amount of buffering and "byte manipulation" capability at the periphery of either system. It is actually the manner in which internal system resources are managed and utilized that provides a useful measure of comparison between them.

Briefly, in a circuit switched network, the source and destination are connected by a dedicated communication path that

is established at the beginning of the connection and broken at the end. This type of connection was specifically selected for use in switched telephony, where subscribers require a continuity in voice transmission and reception. Since the communication path remains fixed for the duration of a conversation, the output speech signal appears to be a time translate of the input speech signal as far as the ear can tell. In addition, for most voice conversations, the allocated analog voice channel is used in a fashion that seems reasonably efficient to the average user.

To establish a connection, the subscriber provides the local central office with an address which is used in setting up a path. Central office equipment detects off-hook, provides a dial-tone, retains dialed digits, generates ringing, busy signals etc. In the current telephone plant, long haul circuits are primarily multiplexed analog channels. Routing selection is performed using a set of prespecified paths and usually based on the first few dialed digits. Call set-up times generally take between 5 and 25 seconds depending upon the number and type of central offices in the link and the amount of traffic. Recent experience has also indicated that reliability and overload problems are becoming increasingly prevalent in certain high density population areas.

A message switching system accepts, transmits, and delivers discrete entities called messages. In such a system, no physical path is set up between the source and the destination and no

resources (e.g. capacity, buffer storage etc.) are allocated to its transmission in advance. Rather, the source includes a destination address at the beginning of each message. The message switching system then uses this address to guide the message through the network to its destination, provides error control, and notifies the sender of its receipt [18].

A simple form of message switching system employing a single central switching computer is commonly referred to as a "star" configuration and has all its lines connected to the central message switch. For many local applications this configuration can be quite practical. Three of its main disadvantages are 1) the central switch may be an unreliable link which will disrupt all communications if it fails, 2) the total circuit mileage for geographically distributed users to connect to the switch may be substantially larger than necessary, resulting in excessive communications cost, and 3) every circuit failure can result in some loss of user communications.

A distributed message switching system is one in which many distributed switching computers are employed and the network control is decentralized in such a way that the failure of any switching computer disrupts communications only for its local customer. The distributed system is usually more economic and reliable than a star configuration for handling geographically distributed users.

The components of a message switching system are dedicated point to point communication circuits and switching nodes which innerconnect the circuits in such a way that a message arriving on one circuit may be transmitted out another. Communication over a message switched system occurs via a sequence of transmitted messages , each consisting of its address followed by text. The address is inspected by each node in routing the message to the next node on the way to its destination. In the ARPANET, one or more computers may be directly connected to a node and are known as Host computers, or Hosts for short. The nodes are called Interface Message Processors or IMPs for short.

A distributed message switched network, such as the ARPANET, contains no mass storage, and as little buffering in the nodes as necessary to utilize the full capacity of the communication circuits*. The network design allows a message to remain in the net only as long as necessary to transport it from source to destination; no long term storage is provided in the communication system. Messages that cannot be delivered to the destination are simply not accepted into the net and must be retransmitted at a later time. Clearly, one or more Hosts on the net with low cost per bit bulk storage could provide or even be dedicated to providing long term storage of messages with subsequent automatic retransmission.

* The Defense Departments Autodin network, however, employs mass storage in the communication network for deferred retrieval and delivery of messages.

The combinatorial aspect of the interconnection of large numbers of computers is an important consideration in network design. Each computer in a message switching system is connected to the net via a single full duplex channel to its IMP over which messages are multiplexed. This single connection to the network makes the computing service accessible to all computers and all users on the net. Furthermore, all users and all computers on other digital networks can access this computer by the simple expedient of a single interconnection between nets. Thus, not only is complete digital access possible, it is achieved in a strikingly economic way for each installation. This technique solves a massive combinatorial access problem with a single economic stroke.

In Figure 1, we show the communications portion of the ARPANET as of April, 1972 when it consisted of twenty-four nodes and twenty-eight circuits. Since that time it has grown to over thirty nodes. Each node is a possible source and destination of messages. We assume (for the moment) that messages may be of variable length up to a maximum of 1000 bits* and are known as packets while in the network. The path taken by a packet traversing the net from node 1 (UCLA) to node 6 (MIT) is indicated by arrows in the figure. The circles indicate the nodal processors and the lines indicate synchronous point to point circuits.

*In the ARPANET, messages may actually vary up to 8195 bits in length.

The message enters the net at node 1, which examines its address and decides to transmit it out its circuit to node 7 (RAND). Upon receipt, node 7 examines the address and decides to send it to node 21 (TINKER), which in turn sends it to node 5 (BBN) which sends it to node 6 (MIT). Node 6, discovering the message is for itself, replaces the destination address by the source address (which is carried along by the message switching system) and "delivers" the message. The text of the message thus appears at the destination exactly as it was transmitted and the address portion identifies the sender. After delivery, the sender is notified of its receipt by a small message that goes back across the network.

An important part of a computer network design is the specification of the location and capacity of all circuits in the net. Fifty kilobit/second circuits are currently used in the ARPANET to achieve an average delay of 0.2 seconds or less. Programs have been developed that iteratively analyse various possible network configurations and select reliable, high throughput, low cost designs through the use of circuit exchange heuristics [20,22]. Analytical techniques have been developed for estimating the average transit time delay under assumed traffic loads. These techniques show that the delay remains almost that of an unloaded net until the capacity of one or more "cutsets" begins to saturate [7].

An important design consideration is the method for

dynamically selecting routes. (We assume that routes are not allocated in advance.) A central controller could provide the routing information and distribute it to all the processors, or the processors could collaborate in computing the routing information directly. This is but one of many instances of a design choice between distributed and centralized control. In the initial ARPANET design, the route selection is performed independently by each IMP according to a distributed routing algorithm. Routing information is stored in a table and individually maintained by each IMP for rapid look-up. It identifies the output line to select for each destination and is updated according to a rule evaluated periodically (e.g. every half second). It could also be evaluated asynchronously (whenever status changes occur) or a combination of both. In the simple algorithm used in Figure 1, each IMP sends the message on its choice of a path with the fewest intermediate IMPs and, using the update procedure, each IMP adapts its routing to other IMP and circuit failures.

A simple method for implementing this algorithm is for each IMP to keep a table with the count of the number of IMPs on the shortest path to each destination which it frequently transmits to its immediate neighbors. Each IMP also announces to its neighbors that it is 0 IMPs away from itself. Upon receipt of the routing information from its neighbors, the IMP increments the neighbors counts and keeps the lowest value for each destination.

Each IMP buffers a packet until receipt is acknowledged by the adjacent IMP. A cyclic checksum, generated in hardware by each IMP, is appended to the transmitted packet for error control. If an error is detected by the hardware at the adjacent IMP, or no buffer space exists, the packet is simply discarded and will shortly be retransmitted by the neighboring IMP when a condition (such as a time-out) occurs and no acknowledgement is received.

The design of an efficient network without mass storage requires that the number of buffers be kept to a minimum, and that they be used so that each IMP is able to use its circuits efficiently and to their maximum capacity. This means that the minimum number of buffers must be at least as large as the number of full sized packets which must be stored from the time one full size packet is transmitted until its acknowledgment returns. This number is determined by the circuit propagation delay, the packet size and the circuit data rates, as well as the total number of circuits. To utilize these buffers efficiently, stored packets must be quickly released upon receipt of their acknowledgment or activated for retransmission, as appropriate, in a timely way.

Each IMP contains only a small amount of buffering for messages in transit and no mass storage, and a flow control strategy is needed to insure that the IMPs do not become "congested" thus preventing useful data from being communicated.

This situation is particularly apparent if the network design allows the source or the destination to temporarily stop the transmission or reception of data and then continue without a loss of messages. This is appropriate to time-sharing computers, and is used in the ARPANET, because it allows occasional delays to occur, for example, while a word is stored in memory, or a procedure is activated by the processor. In practice, a Host can neither guarantee to accept all messages at their instantaneous arrival rate, guarantee not to crash while receiving heavy traffic, or expect the transmitting Host to buffer messages should he prefer to discard them upon receipt. In particular, flow control is necessary to protect the network against the sudden dispatch of a larger number of messages to a single destination than it is prepared to accept [13].

An often overlooked but important consideration in the network design is whether or not to keep the circuits fully loaded even in the absence of maximum traffic. For instance, should "test messages" be continuously transmitted or only periodically transmitted to check circuits? Under light traffic loads, is it desirable to transmit duplicate packets and accept the first one with a valid checksum, in order to reduce occasional retransmission delays or to improve the response time on a very noisy circuit. For land based circuits, the extra traffic during otherwise light loads appears to be acceptable and desirable to reduce delays on noisy circuits. The extra processor capacity is ordinarily available for heavy traffic in

any event. For multi-access satellite circuits, however, the extra traffic during light loads may interfere with other processors sharing the same channel.

In an unloaded net, the transit time is determined primarily by the number of IMPs encountered in the routing and the time for the packet to pass from one IMP to the next. This, in turn, is determined directly from the packet length, the circuit data rate, and the speed of light propagation delay. Under increasing traffic loads, the transit time also begins to increase due to occasional delays in the IMPs. However, if the network is designed to begin rejecting the further input of traffic as the capacity limitation of the network is approached, these delays can be kept to a few times that of an unloaded net. Traffic is thus queued outside the net (rather than allowed to enter and be queued inside the net) so the nominal transit time during peak traffic is not very different from that experienced in an unloaded net. In these cases, an attempt must be made to insure that the effective bandwidth is shared "fairly" among all the competing sites.

Network usage generally requires a combination of short transit times for interactive usage and high bandwidth for file transmission. These two objectives may be attained with single packet messages. To achieve interactive transit times, no set-up delay must be incurred. A simple way to achieve this is for the source IMP to retain a copy of each packet which is nominally

discarded after the delivery is made, but retransmitted when, for lack of buffer space, the original is discarded at the destination IMP. To achieve high bandwidth, enough messages must be allowed to enter the net between source and destination so as to fill the "pipeline", but this flow must be able to be readily quenched at the source when the buffer space at the destination IMP begins to fill.

The current ARPANET design actually allows variable length messages with a maximum size just over 8000 bits. The message is partitioned by the source IMP into separate 1000 bit packets to speed its transmission through the network. Each packet makes its way to the destination independently where it may conceivably arrive out of order. These packets could be reassembled into the proper order by the destination Host (using sequence numbers), but when the assumption is made that the communications net should preserve sequencing at least at the level of a single process to process conversation, the IMPs are obliged to reassemble the packets. The destination Host thus receives the text of each message exactly as it was transmitted in a single block transfer.

When these larger messages are used and the IMPs undertake the responsibility for reassembly, yet another type of congestion phenomenon called reassembly lockup is introduced [13]. The flow control mechanism which is used to prevent the backup of messages in the net is also powerful enough to prevent the lockup problem.

But in its application, it can subject long messages to set-up delays and thus delay succeeding short messages from its Host.

If "sufficient" buffer space were available for reassembly at the destination IMP there would be no "a priori" compelling reason for the subnet to preclude a Host from sending full 8000 bit messages (or even somewhat larger ones). However, the presence of 8000 bit messages may noticeably delay shorter messages from other Hosts 1) while it is being delivered to the destination and 2) by tying up eight buffers rather than one during reassembly. This provides one valid reason to restrict the Hosts to single packet messages, if these delays become significant. However, as we indicate in section V, there may be other factors which favor retention of the larger size.

If there is a fundamental distinction between circuit switching and message switching, it is undoubtedly in the way internal resources are managed. For example, circuit switching requires that network bandwidth as well as local control equipment and storage be allocated to a given transmission circuit in advance, whereas a message switching system does not. Secondly, the presence of circuit switched routes means that user messages are identified by their circuit and no user control signalling need accompany the transfer of information. In message switching, however, all record of activity (except accounting) associated with a message is contained in the message, which vanishes when the message leaves the system. This

signalling information, in the form of an address, must accompany each message and the message must be examined and processed at each stage of the transmission process.

Two practical consequences of the difference are that the circuit switched system usually requires a non-trivial set-up time to allocate resources. Message switched systems can avoid set-up delays, but may introduce occasional variations in transit time. These delays can generally be maintained to within a few times the delay of an unloaded net, but wider variations may result from queueing delays outside the net, particularly under heavy traffic load. Under similar conditions, though, a circuit switched user might fail to obtain a circuit and would incur this probabilistic situation on subsequent tries. Any allocated but idle channels are simply unavailable at this time to handle these overload conditions.

V. NETWORK USE OF INDIVIDUAL COMPUTER SYSTEMS

The term network has been used and misused in a variety of ways. Some people have referred to the use of dial-up facilities to access a single computer as a network. Others have referred to any interaction between computers and a communication system as a network capability. Several distributed networks were developed to allow simple forms of communication between identical machines using standard dial-up or leased voice circuits, thus providing a convenient way to transfer jobs and files and to maintain and update the systems in the net. This latter application exemplifies a true networking activity, even though it only concentrates on selected aspects of computer resource sharing [19,21].

In general, only a subset of the network sites possess computing power, and certain of them will offer regular service to users via the network. Other sites may choose to offer service only on a limited basis, or to cooperatively interact via the net with selected co-workers, but not offer general service. This latter situation is more likely to occur for many specialized research facilities. In addition, large private computer centers as well as commercial firms may welcome the opportunity to connect their systems since it offers a large potential market for usage of unused capacity.

Host service on the network ought to be as reliable as the communications, although this objective is often difficult to

achieve. For example, in the ARPANET, total uptime of the IMP at any site is currently on the order of 98 or 99%, while Host availability is generally no higher than 90%. It is certainly possible to improve on this score; some commercial firms claim to provide over 99% availability of service, and certain private and government systems must obviously be operated with near perfect reliability. The airline systems and the computers in the space program provide two key examples.

a. Network Access

In a network that supports computer to computer communication, user groups with a local computer can access another computer in the net by first logging into their home computer and then into the other computer, using the home computer in a transparent mode as a switch. However, this is an expensive way to access another computer since it incurs charges in two computers and ties up jobs in both. Furthermore, since Hosts may be unreliable, the connection is more vulnerable than a direct connection into the other machine.

Sites with expensive computer installations might deem it economic to depart from their machine if "equivalent" service can be economically obtained via the net. In general, such a site requires the ability to service local users with a full complement of terminals and peripheral devices such as teletypes,

graphics displays , line printers, magnetic tapes, and possibly other mass storage. In addition, many sites with no computer will derive maximum benefit in connecting to a net if the full range of peripherals can be provided locally. Other users, however, may be satisfied with a simpler approach that provides terminal access to remote computers but does not handle peripheral devices other than line printers.

For this latter class of users, an addition may be made to the IMP which allows a set of terminals to be directly connected. This addition consists of multiplexing equipment that collects characters from the terminals and packages them in the form suitable for delivery to the IMP. Likewise, it receives messages from the IMP and sorts the characters out to the various terminals. This addition requires hardware and software designed to make the set of terminals interface to the network as a "mini-host", and this IMP is referred to as a Terminal IMP.

A more elaborate approach is appropriate for user sites that wish to support many different types of peripheral equipment. Since their characteristics and operation can vary widely, device dependent programming is mandatory, and substantial buffering may be required for the higher speed devices. Furthermore, sites with mass storage will generally wish it to be accessible from other computers in the net, which generally requires the implementation of a full set of standard and specialized network protocols. These considerations make it appropriate to provide a

separate processor devoted to the handling of peripherals.

This latter approach is particularly desirable for a site which is both a user site and a server site. The architecture of such a site should be organized so that if the serving site cpu is down, local users can access other network computers and local mass storage can be accessed by them over the net. Similarly, if the local storage should fail, others across the net can temporarily replace them. If the net fails, local users can still obtain full access to the local system. Only if both the local cpu and the network fail will the users be unable to obtain computation.

Modularity and logical reconfiguration are conveniently achievable in this way. Substantial progress in the design of modular communication oriented architectures can be expected from innovative usage of interconnection ideas [3,4,5].

b. User Requirements

Let us now turn our attention to the use of these facilities by the user. We note three potential locations where user computation can be performed - in the terminal itself, in the peripheral processor (or Terminal IMP) or in the Host computer(s). Although the bulk of the "pure computation" will undoubtedly take place in the Host computers, some aspects of the processing must, in general, be distributed. For example, local

echoing is required to obviate the otherwise noticeable effects of speed of light propagation delay as on satellite links. This raises the important question of location of functions in a distributed network. In other words, what intelligence is needed to allow distributed system usage and where should it be placed?

Let us concentrate on the echoing problem for the moment. As a general rule, a remote user should see the same output and otherwise obtain service from a remote Host as if he were a local user. To achieve this objective, the users local system (programmable terminal, "mini-host", etc.) must have considerable information available to it about each subsystem in use at the remote Host. For example, a simple local echo/remote echo strategy is generally insufficient to handle echoing for users on half duplex terminals, or users on full duplex terminals that prefer to type ahead. A remote user editing the character string ABCDEFG can delete the last three characters by typing successive delete characters (echoed as) and he sees the output ABCDEFG G F E. Using the system from a remote site with local echoing and typeahead he would see ABCDEFG GFE.

More striking, perhaps, is the remote use of a debugging program DDT. To examine successive registers 120 - 123 a local user would first type 120/ to print the contents of the first register and then strike successive line feed characters to examine the successive registers. The system would respond with (say the contents are all zeroes):

120/ 0

121/ 0

122/ 0

123/ 0

and leave the cursor following the last 0. A remote user with typeahead and local echo would see one of several possible responses (depending on the remote systems response to a received line feed). Assuming the remote system echoes only the formatted data and the local system echoes linefeed as linefeed (no carriage return), the output would look as follows:

120/

0121/ 0122/ 0123/ 0

In general, the local echoing system should have full knowledge of the time varying syntactical operation of the subsystem in use. This requires feedback of information about subsystem break and separator characters, control signalling, special conventions, etc.

Each terminal has characteristics peculiar to it and a convention is required for a computer system to initially recognize a terminal. Although the remote computer could then convert to the characteristics of the terminal, it is far more manageable if each terminal could appear to the network as a standard terminal employing an agreed upon set of characters and signalling conventions. One such standard (developed for the ARPANET) is 7 bit USASCII with the eighth bit set to 0. In that scheme, the other 128 possible characters are reserved for special control characters. In addition, most terminals need

attention to details such as carriage return, keyboard locking and unlocking, interrupt signalling, and other peripherals indicate out of paper, buffer full, and may require complete two-way channel control, etc.

Local computation is therefore needed for the users terminal to interact properly with other remote systems and their subsystems. At a minimum, his local computation must allow the user to 1) identify his terminal to the network, 2) select a destination Host, 3) select a transmission mode, 4) perform echoing and code conversion, and 5) allow the remote Host to be interrupted.

It seems probable that, in the long run, many terminals will contain mini-processors and thus user programming in a separate "mini-Host" will be unnecessary. However, until this possibility is a widespread reality rather than an expectation, users may be hindered if they are unable to provide local user code in one place or the other.

c. Message Processing

Before considering various examples of usage of a computer network, let us briefly indicate how messages are processed within the Host computers [15,32]. Messages travel through several layers of protocol in the Host system. The first layer of protocol handles the IMP, activating I/O buffers, fielding

control messages, etc. The second layer interacts with the local processes and remote Hosts monitor, allocating buffer storage, providing process identification, formatting control information etc. Subsequent layers correspond to specific user oriented functions such as the standard network terminal, file transmission, etc.

The ARPANET Host protocol utilizes the notion of connections over which messages are transmitted. A connection must first be established before communication over it may occur. The Hosts at either end of the connection must keep full information about the use of the connection (which is obtained during its establishment) to handle flow control. This strategy appears to be close in spirit to telephone circuit switching.

A few limitations to this strategy are apparent [14]. An important concern is that it requires each Host to maintain resources in the form of connection tables that can become filled thus preventing any further communication with that host. In particular, a single process can attempt to establish its maximum limit of connections although it cannot, in general, make full use of them at one time. Entries in this connection table are permanently allocated and thus only a fixed number of connections can be established at any time.

A second limitation with the use of a connection table is that it can be vulnerable to error conditions and Host status, particularly since both Hosts must generally agree on its

contents for flow control. Finally the strategy requires the connection information to be used for termination which means that information which otherwise would be nominally discarded by Hosts with limited space, must be retained merely to close the connection. These limitations , as well as others, may be obviated with a message switched Host protocol [14].

The desired size of Host messages may have an important impact on the operating system as well as on the communication system. The original ARPANET design specification allowed individual messages to be as large as 8192 bits, a design choice based largely on intuition.

As the design specification originally stated [29]:

"... a packet is defined as the inter-I²P unit and Message as the inter-Host unit. A packet will not exceed 1024 bits in length. The I²Ps must break all longer messages into multiple packets. Messages will be limited to 8192 bits so as not to require excessive buffer space."

Undoubtedly, this latter reference is interpreted as referring to buffer space in the I²Ps, but it could equally well apply to buffer space in the Hosts. In particular, the argument defends why the size is not larger but does not entertain the possibility that it ought to be kept smaller for any specific reason.

If there is a convenient maximum Host message size, it is probably a maximum sized page, which corresponds to 1K of 36 bit

(or possibly 48 bit) words. However, transmission of such large messages (say 50,000 bits) to the IMP and from the IMP to the Host produces excessive delays for short messages queued up behind them, and provides a prime reason for Hosts to prefer that these long messages be subdivided into shorter messages. Since no experience with network software was available during the initial design, it was intuitively concluded that a shorter 8192 bit message was short enough. Interestingly, we note that two hardware paths between each Host and its IMP, one for short messages and one for long messages, could remedy this problem at some extra cost in hardware and buffer storage.

Since efficient transmission is possible with 1024 bit packets, it appears in retrospect that the selection of the larger message size may be unnecessary. The Host overhead in network communication increases with the number of messages, so there is some incentive for making all Host messages sufficiently large that a typical short transaction can occur in a single message. No evidence yet obtained by us indicates that 1024 bit Host messages would impose a limitation that is significant, but an increased demand for page transfers or the presence of higher bandwidth circuits could tip the balance more strongly in favor of a larger size. The jury seems to be still out on this issue.

VI. APPLICATIONS OF MULTI-COMPUTER INTERACTION

Network utilization involving the combined use of two or more computers in a productive way began during the initial experimentation with the ARPANET. It has provided experience in the development of techniques for performing distributed computation and allowed some simple application areas to be identified. Some applications involving multiple computers have been discussed for many years, partly as a result of their inherent interest and ease of conceptualization. One important example is the access to specialized data bases that are only available from a remote source. Several information banks have already been developed or are under development, and their expected usage is being projected upward. Another example is in the use of "future" computer communication networks for handling the distribution and delivery of mail and other transient information. However, these applications are only beginning to develop in any significant way. Much effort has already been devoted to the study of topics such as concurrency and parallel processing which may result in faster program execution and otherwise make efficient usage of available resources. We expect that computer networking will enhance these efforts. For other applications, the sensation of dealing with one system rather than two (or more) is overwhelmingly evident to the user and this pleasant feeling often generalizes to other multi-computer interactions as well.

Three areas in which applications have already occurred are briefly identified below:

i. File Transfer

The first application for combining two computer systems in the ARPANET in a non-trivial way involved the use of an XDS-940 computer at SRI and a PDP-10 computer at UTAH. SRI, anticipating the delivery of a PDP-10, began to use the UTAH machine in the development of PDP-10 software.

At first, a higher level language was developed. Source code was generated on the 940, converted to object code and executed on the 10. Patches were made on the UTAH machine during debugging and, periodically, an updated source and binary version would be generated at SRI and sent over the net. Subsequently, other higher level languages were similarly developed.

A simple protocol to handle file transfers was developed for the TENEX operating system [10] and has proven useful for transferring new subsystems and system revisions between TENEX sites. In addition, it has been a useful initial step to allow cooperating processes in two TENEX systems to share a single file.

In this protocol, the network appears as a device to which a file may be output or from which a file may be input. The two

ends of the transfer must coordinate by having one end execute the input and the other end execute the output. This simple file transfer protocol requires the intervention of the user to log into both ends, assign a file name for the destination, invoke the proper format, etc. Other experience in the transfer of files has been recorded by the University of California at Santa Barbara, as well as by IBM, by Control Data and others [19,21].

ii. Remote Job Service

A simple example of a computer to computer interaction is provided by users who write, debug, edit and store programs on an interactive time-sharing facility and run them on a separate batch processing system. While time-sharing has created an interactive environment for programming and the development of programming techniques, batch processing systems (and small dedicated computers) have maintained a predominance for performing extensive computations. The availability of both kinds of service in a computer communications network provides a single user with convenient access to the best features of both. See figure 3.

While a user can become accustomed to using both services independently, he need not be required to physically collect time sharing output on tape or punched cards for submission to the batch system. The most convenient user option is for the

interactive machine to submit his job to the batch processing machine under user initiated control. He can then specify the location for output to be stored or printed, revise the program in the time sharing system and resubmit it under fully interactive procedures from a single location, with no need to keep physical copies of files, program etc.

For over a year, the RAND Corporation had been using the ARPANET for remote job service from an IBM 360/65 at RAND to a 360/91 at UCLA [26,27]. This facility was only accessible to internal RAND users until recently when it was replaced by a PDP-10, which allowed network users to create and submit jobs for remote service elsewhere.

Both the program and its relevant files must be transferred to the 360/91 before a job can be run. They are typically shipped together as successive "card images". The remote job service program will allow the users to start or stop the execution of his program, cause the system output to be stored on a designated file, or be output on a device such as a local printer. The user is also provided with options to check the status of the execution, receive confirmation and error messages that indicate its progress, and allow certain actions to be taken.

These facilities are used by RAND researchers in the generation and processing of simulated weather data. Weather

modelling programs can be activated from a remote site, output from these programs can be temporarily stored or shipped to a remote site for preparation and display. This separation of the computation into components is particularly appropriate when one part may be devoted almost exclusively to extensive numerical or symbolic computation and another part to user related manipulation of or preparation of the output data.

A "complex" weather simulation program requires many hours of computation on the 360/91 and thus is not well suited to rapid on-line activities such as the updating of a display. Rather, pre-computed weather data (from the models) is retrieved from 360/91 disc packs (with operator assistance) and used by the PDP-10 for further processing and display. The availability of a high speed parallel processing system such as the ILLIAC IV [33] may eventually allow real-time weather experimentation without operator intervention.

iii. Multi-Process Operation in Many machines

The combined use of two or more computers allows additional processing capability over the use of a single system. One such example is provided by the McRoss system [12] that coordinates the operation of two or more cooperating air traffic control simulation programs running in one or more TENEX systems. Each simulation program, called ROSS for Route Oriented Simulation

System [23], models the airspace of one air traffic control center in detail. To simulate the airspace of a Boston to New York flight, four simulation programs would be activated; one for the Boston terminal area, one for the Boston enroute area, one for the New York enroute area, and one for the New York terminal area. The four Ross programs may be run simultaneously in as many as four TENEX systems in the ARPANET.

When a single machine is used to house all the components of a programming system, it has the disadvantage that computing will stop if that machine crashes. When one piece of a multi-computer programming system becomes unavailable, the other parts can learn to adjust to the change in configuration. A desirable objective is to provide enough backup information to enable the multi-computer programming system to be restarted in the event of a single Host failure and to proceed from a recent point in simulated time as if nothing had happened.

Other applications involving multiple computers are certain to arise for which simple examples are more difficult to construct. For example, as special areas of expertise develop, it is natural to expect that individual efforts by specialists also trained in the use of computers, will produce new and useful resources on different machines. These resources may represent state of the art or proprietary developments that cannot be conveniently transferred to other machines and must therefore be used at the site of their creation or where they currently exist.

An important application for distributed computation is thus likely to involve the coordination of separate research projects into combined efforts that utilize these specialized or proprietary and hence non-transferable resources.

A second major application of distributed computation is likely to be the facilitation of interactive cooperation between people at different locations. Interactive cooperation may be regarded as an extension of normal voice communication to include the ability for several persons at different locations to "simultaneously" observe, communicate about, and manipulate both common data structures and programs. Since the people are assumed not to be co-located, the programs which support the interactive cooperation (such as display protocol routines) must also be distributed.

A third major application for distributed computation is in providing for conveniently feasible demonstrations of prototype systems to be performed from different locations. This technique can allow new capabilities to be readily conveyed without the inconvenience of moving the observer to a home site for the demonstration.

VII. A DISTRIBUTED OPERATING SYSTEM

A network in which basic differences exist between the computers at each installation is said to be inhomogeneous. It is possible to develop a standard network protocol for an inhomogeneous system that allows usage of various pieces of the system to be coordinated in a uniform manner. However, this task is one of substantial complexity that will probably require changes in system architecture and program design techniques before it can be fully realized. Even if it were a straightforward matter, it would not be generally useful to transfer portions of any one system to another, and standard operations that involve systems at a remote site must typically be performed at that remote site.

A collection of similar operating systems may also be organized into a virtual subnetwork of homogeneous computers that interact with each other in a uniform way. These systems are more easily organized into a single distributed operating system with common file systems, address space, naming conventions etc. In general, every type of interaction between two systems in a homogeneous network must be evaluated to determine what is to be transmitted and what is to be remotely evaluated. No single answer will suffice for all applications. As we noted above, it is not generally possible to apply both alternatives in an inhomogeneous network. We consider some of the properties a system like this ought to possess.

The user accesses the distributed homogeneous network by logging into a distributed system rather than into a specific computer in the net. An appropriate machine is selected for him and he logs in with the standard login sequence for his home computer, including password, account number and other information as required. Upon completion of the login sequence, the computer initiates a brief exchange with the users home computer to notify it of the impending job which it then proceeds to service. The home computer may then request that the job be transferred, alter credit or accounting information, or merely note the event.

Under conventional design constraints, the combined operation of several Host computers will require a separate job to be established in each machine. In a distributed system, though, it is important to allow access to each system without the user logging into each system individually. Furthermore, it is also desirable to permit certain transient activities to occur, which do not tie up valuable resources or otherwise interfere with users on the system. The system merely performs the transient activity and logs the transaction into a suitable file for accounting purposes. Once logged into one of the Hosts, the user is able to access and utilize any programs, files and most other facilities on other computers in the system as if they were all on one virtual machine.

The availability of many resources in this system makes it

possible to achieve reliable operation when one or more resources are disabled. The user can be affected by failure in several ways, for example, his program or a piece of it may be aborted by machine failure or he may lose part or all of his files. The user may also find the local file storage to be unusable while running his job. If local storage is not initially available, he can specifically designate another system to store his files. Alternatively, he can allow the local system to store files in other Hosts and expect them to be returned without his knowing the identity of the temporary storage location. Obviously, a small amount of local storage is needed for this application. The distributed system thus not only makes resources more available, it can use them to provide increased reliability to a user.

A system designed to operate stand alone may not be as efficient in serving its network users as in serving its local users. Certain performance improvements are obtainable by streamlining of critical portions of the system code, attention to organizational details and to carefully engineered improvements to scheduling, the file system etc. However, a major improvement in speed and efficiency may require structural overhaul of the system organization to allow for efficient process to process communication at high bandwidths, and for efficient overall utilization of resources. In particular, the portion of the system devoted to protocol and message handling

(byte manipulation) can consume a considerable amount of cpu time at high bandwidth. In a time sharing system, particularly, these functions, which appear to be communication functions, can be usefully separated from the "computation" functions. This decomposition also allows more human engineered network interfacing since, a local cpu failure can usually be reported by the communications portion of the system if it remains up. The delegation of all protocol functions to a separate processor that can directly deposit into and retrieve from process buffers makes it possible for the operating system to communicate over the network at speeds at least an order of magnitude faster than before. This increase in capability is achieved by performing the protocol operations on the fly in a separate processor and by avoiding the unnecessary overhead in moving real-time data around in memory.

Techniques for computer to computer communication are still in their infancy and a great deal of exploration and experimentation is occurring in this area. How should programs be written to run in a network environment and what debugging and control techniques are suitable for distributed computation? What operating system architectures are appropriate to computer communication? The efficient utilization of a distributed operating system involves the sensible decomposition of a task into components. This requires timely access to status information and the ability to use this information wisely in the allocation of tasks to resources and in their scheduling. Just

as the management of communication resources was central to the operation of a communication subnet, so will the management of computer resources be to the overall utilization of a computer network.

VIII. CONCLUSIONS

A principal motive underlying computer network development is to provide a convenient and economic method for a wide variety of resources to be made available and to be shared. Such a network provides more than an increased collection of hardware and software resources; it affords the capability for computers as well as individuals to interact in the exchange and processing of information.

It is not usually the case that a program written for one computer can be shipped to another computer and run there to completion, correctly. It may be possible in a number of cases where the machines are nominally identical, but it is usually the case that a program must be run on the machine for which it was written. It is thus desirable to strive for compatibility between at least a subset of the system resources, including the use of machine independent higher level languages, the use of network wide standard protocols, or the use of nominally identical systems.

The development of communication subnets has been strongly influenced by the regulatory climate and the need for reliable and economic ways to achieve both remote terminal access and high bandwidth switched computer to computer communication. Message switching has emerged as a strong contender for computer to computer communications. It has been demonstrated to provide a highly reliable, error-free method of achieving interactive

switched communications. Although its technical feasibility has been firmly established, its practical utility is under evaluation, and under close scrutiny, it may prove to be a viable economic alternative to conventional circuit switching.

It is important that a communication system not preclude the possibility that separate or private data networks may be accessed through it in a standard and convenient way. A digital message switched network has this property while an analog frequency based system may not. Incompatible data networks are clearly undesirable if all resources are to be mutually accessible. If separate data networks are jointly planned before development, at least at the interconnection level, they may be connected at a later date and viewed together as a single network that evolved by way of separate networks.

The great diversity of resources in a computer network may initially hinder its growth. Users must familiarize themselves with many different systems often without the aid of substantive interaction with systems personnel or clear and complete documentation. But the potential benefits of computer networks are sufficiently great that, over time, this obstacle will surely be surmounted and in the process may lead to superior standards for system operation and documentation.

Computer networks provide a unique mechanism for increased participation between individuals. Participation in research and development using the distributed resources of a computer network

can lead to the close cooperation between individuals who might otherwise have little incentive to work together. This interaction can further cross fertilize the network community and encourage even higher levels of achievement through technical cooperation.

REFERENCES

1. "Pinched Budgets Promote Growth of University Computer Networks," Communications of the ACM, Vol 15, No. 3, pp.206-207
2. P.M. Walker and S.L. Mathison, "Regulatory Policy and Future Data Transmission Services," Chapter Nine of Computer Communication Networks, edited by Abramson and Kuo, Prentice Hall, 1973
3. D.Farber, K.Larson, The Architecture of a Distributed Computer System- An informal Description, University of California, Irvine, Technical Report NO. 11, Dept of information and Computer Science, 1970
4. H. Baskin, B. Borgerson, R. Roberts, PRIME - An Architecture for Terminal Oriented Systems, AFIPS SJCC, 1972, pp. 431-437.
5. R. Davis, S. Zucker, C. Campbell, A Building block Approach to Multiprocessing, AFIPS SJCC, 1972, pp. 685-703. Burroughs, the interlock
6. S. Crocker et al, Function Oriented Protocols for the ARPA network, AFIPS SJCC, 1972, pp. 271-279
7. H. Frank, R .E. Kahn, L. Kleinrock, Computer Communication Network Design- Experience With Theory and Practice, AFIPS SJCC, 1972, pp. 255-270
8. L. G. Roberts, B. Wessler, The ARPA network, In Computer Communication Networks, edited by Abramson and Kuo, Prentice Hall, 1973

9. M. Beebe, N. Sullivan, Tymnet - A serendipitous evolution, Second ACM IEEE Conf. on Problems in the Optimization of Data Communication Systems, Palo Calif. 1971
10. D.G.Bobrow, J.Burchfiel, D.Murphy, R.Tomlinson, TENEX - A paged Time-sharing System for the PDP-10, Comm. of the ACM, Vol 15, No. 3, March 1972, pp. 135-143
11. F. J. Corbato et al, An Introduction and Overview of the Multics System, AFIPS FJCC 1965, pp. 185-196
12. J. Thomas, D. Henderson, McRoss - A Multi-Computer Programming System, AFIPS SJCC 1972, pp.281-293
13. R.E.Kahn, W.R. Crowther, Flow Control in a Resource Sharing Computer Network, IEEE trans on Comm. Technology, Vol. COM-20, June 1972, pp. 539-546
14. D.C.Walden, A system for Interprocess Communication in a Resource Sharing Computer Network, Comm. of the ACM, Vol 15. No. 4, April 1972, pp. 221-230
15. The Host/Host Protocol for the ARPA Network, Network Information Center, Document No. 147, Stanford Research institute, Menlo Park, Cal.
16. L.G.Roberts, B.Wessler, Computer Network Development to Achieve Resource Sharing, AFIPS SJCC, 1970, pp. 543-549
17. P.Baran, S. Boehm, P. Smith, On Distributed Communications, Series of Rand Reports, Rand Corporation Santa Monica Calif, 1964
18. F.Heart, R.E.Kahn, S. Ornstein, W.C. Crowther, D. Walder., The Interface Message Processor for the ARPA Computer Network, AFIPS

SJCC, 1970, pp. 567

19. A.H. Wels, Distributed Network Activity at IBM, IBM Research Report RC3392, June 1971

20. H. Frank, I. Frisch, W. Chou, Topological Considerations in the Design of the ARPA Computer Network, AFIPS SJCC, 1970, pp. 581-587

21. W.J. Luther, Conceptual Bases of CYBERNET, Courant Symposium no. 3, 1970, Reprinted in Computer Networks, Prentice Hall, Englewood Cliffs, pp. 111-146

22. R. Van Slyke, H. Frank, Reliability of Computer Communication Networks, Proc. of Fifth Conf. on Applications of Simulation, New York, Dec. 1971

23. W. Sutherland, T. Myer, D. Henderson, E. Thomas, Ross, A Route Oriented Simulation System, Proc. of Fifth Conf. on Applications of Simulation, New York, Dec. 1971

24. T. Marill, L.G. Roberts. Toward a Cooperative Network of Time-Shared Computers, FJCC 1966

25. Project MAC Conference on Concurrent Systems and Parallel Computation, Woods Hole, Mass., June 1970

26. E. Harslem, unpublished memo, Rand Corporation, 1972

27. R. Braden, unpublished memo, University of California at Los Angeles, 1972

28. C.R. Fisher, Introduction to the DATRAN Switched Digital Network, Proceedings of the International Conference on Communication, June 1971, pp. 23-1 to 23-3.

29. Request for Quotation (Interface Message Processors for the

ARPA Computer Network), Defense Supply Service, Washington, D.C.,
July 1968


30. R.A.Scantlebury, P.T. Wilkinson, The Design of a Switching
System to Allow Remote access to computer Services by other
Computers, Second ACM IEEE symposium on Problems in the
Optimization of Data Communication Systems, Palo Alto,
California, October 1971

31. D.W. Davies, K.A. Bartlett, R.A. Scantlebury, and P.T.
Wilkinson, A Digital Communication Network for Computers giving
Rapid Response at Remote Terminals, ACM symposium on Operating
System Principles, Gatlinburg, 1967


32. Bolt Beranek and Newman Report No. 1822, Specifications for
the Interconnection of a Host and an IMP, Cambridge, Mass.
02138, October 1971

33. W.J.Bouknight et al., The Illiac IV System, Proc. Of the
I.E.E.E., Vol 60, No. 4, April 1972, pp. 369-388

34. Datamation, April 1972, pp. 36-46



24-IMP NETWORK - 4/1/72



TERMINAL IMP

A SIMPLE IMP NETWORK

